Glide 5.5

Quick Start Guide

Schrödinger Press

June 2009

Contents

Document Conventions

In addition to the use of italics for names of documents, the font conventions that are used in this document are summarized in the table below.

Font	Example	Use
Sans serif	Project Table	Names of GUI features, such as panels, menus, menu items, buttons, and labels
Monospace	`$SCHRODINGER/maestro`	File names, directory names, commands, environment variables, and screen output
Italic	*filename*	Text that the user must replace with a value
Sans serif uppercase	CTRL+H	Keyboard keys

Links to other locations in the current document or to other PDF documents are colored like this: Document Conventions.

In descriptions of command syntax, the following UNIX conventions are used: braces `{ }` enclose a choice of required items, square brackets `[ ]` enclose optional items, and the bar symbol | separates items in a list from which one item must be chosen. Lines of command syntax that wrap should be interpreted as a single command.

File name, path, and environment variable syntax is generally given with the UNIX conventions. To obtain the Windows conventions, replace the forward slash / with the backslash \ in path or directory names, and replace the $ at the beginning of an environment variable with a % at each end. For example, `$SCHRODINGER/maestro` becomes `%SCHRODINGER%\maestro`.

In this document, to *type* text means to type the required text in the specified location, and to *enter* text means to type the required text, then press the ENTER key.

References to literature sources are given in square brackets, like this: [10].

Getting Started

This manual contains tutorials designed to help you quickly become familiar with the functionality of Glide, using the Maestro interface. This chapter contains a brief overview of the software and some setup instructions for the tutorials. The tutorial begins in Chapter 2 with the generation of grids from a prepared protein to represent the receptor for docking. In Chapter 3, a set of ligands is docked and scored, and the receptor and ligand poses are examined in Chapter 4. Protein preparation is not covered in this manual: see the *Protein Preparation Guide* for details of this task.

Panel-specific online help is available for all Glide panels. If you need help with a Glide task, click the Help button or see the *Glide User Manual*.

To complete the exercises, you must have access to an installed version of Maestro 9.0 and Glide 5.5. For installation instructions, see the *Installation Guide*.

Exercises in some chapters produce structure files that are needed in subsequent exercises. To allow you to begin at any exercise you choose, these and other necessary files (ligand files, for example) are included with the Glide distribution.

1.1 About Glide and Maestro

Glide is designed to assist you in high-throughput screening of potential ligands based on binding mode and affinity for a given receptor molecule. You can compare ligand scores with those of other test ligands, or compare ligand geometries with those of a reference ligand. Additionally, you can use Glide to generate one or more plausible binding modes for a newly designed ligand. Once you locate favorable structures or bonding conformations with Glide, you can use Liaison or QSite to obtain binding energies for ligand-receptor pairs.

Protein Preparation is usually required for Glide calculations. It can be performed for most protein and protein-ligand complex PDB structures using the Protein Preparation Wizard panel in Maestro. For detailed information, see the *Protein Preparation Guide*.

Maestro is Schrödinger's powerful, unified, multi-platform graphical user interface (GUI). It is designed to simplify modeling tasks, such as molecule building and data analysis, and also to facilitate the set up and submission of jobs to Schrödinger's computational programs. The main Maestro features include a project-based data management facility, a scripting language for automating large or repetitive tasks, a wide range of useful display options, a comprehen-

sive molecular builder, and surfacing and entry plotting facilities. For more detailed information about the Maestro interface, see the *Maestro Overview*, the Maestro online help, or the *Maestro User Manual*.

Maestro comes with automatic context-sensitive help (Auto-Help), Balloon Help (tooltips), an online help facility, and a user manual. For more information on getting help, see page 35. You can also undo some operations in Maestro. For more information, see page 31 of the *Maestro Overview*.

The **Impact** computational engine is the underlying computational program for Glide. It can perform molecular mechanics calculations, either through the Maestro interface or from the command line. For information on running basic Impact jobs, see the *Impact User Manual* or the *Impact Command Reference Manual*.

1.2 Preparing a Working Directory

Before you begin the tutorial you need to create a working directory to keep all your input and output files, and then make a copy of the tutorial files.

UNIX:

1. Set the `SCHRODINGER` environment variable to the directory in which Maestro and Glide are installed:

   csh/tcsh: `setenv SCHRODINGER` *installation_path*
   sh/bash/ksh: `export SCHRODINGER=`*installation_path*

2. Change to a directory in which you have write permission.

   `cd` *mydir*

3. Create a directory by entering the command:

   `mkdir` *directory-name*

4. Copy the structure files to the `structures` subdirectory (*version* is the 5-digit Glide version number):

   `cp -r $SCHRODINGER/impact-v`*version*`/tutorial/structures .`

   This command creates the subdirectory as well as copies the files.

5. In the working directory, create subdirectories named `glide` and `grids`.

   `cd` *directory-name*
   `mkdir glide grids`

Windows:

1. Open the folder in which you want to create the folder that serves as your working directory.

   The default working directory used by Maestro is your user profile, which is usually set to `C:\Documents and Settings\`*username* on XP and `C:\Users\`*username* on Vista. To open this folder, do the following:

   a. From the Start menu, choose Run.

   b. Enter `%USERPROFILE%` in the Open text box and click OK.

2. Under File and Folder Tasks, click Make a new folder.

   You can also choose File > Folder > New.

3. Enter a name for the folder.

   If you want to create a folder inside this folder, open the folder and repeat steps 2 and 3.

4. Open the folder that contains the tutorial files. This folder is in the Schrödinger software installation, which by default is installed at `C:\Schrodinger2009`.

   a. Open an explorer window.

   b. Navigate to the Schrödinger software installation.

   c. Open the `impact-v`*version* folder (*version* is the 5-digit `Impact` version number), then open the `tutorial` folder inside that folder.

5. Drag the `structures` folder to the folder you created in Step 3.

   You can close the `tutorial` folder now.

6. Create folders named `glide` and `grids` in the working folder.

   You can use Step 2 and Step 3 for this task.

1.3 Starting Maestro and Setting the Working Directory

Once you have created the working directory you can start Maestro, and set the Maestro working directory. By default, Maestro writes job files to its working directory. You can change the default in the Preferences panel. If you have changed the default, you should change it back for this tutorial.

UNIX:

If you have followed the directions in the previous section, the SCHRODINGER environment variable should be already set, and you should be in the working directory. You can then skip the first two steps.

1. Set the SCHRODINGER environment variable to the installation directory:

   csh/tcsh: `setenv SCHRODINGER` *installation_path*

   bash/ksh: `export SCHRODINGER=`*installation_path*

   This environment variable is also required to run Glide jobs.

2. Change to the desired working directory:

   `cd` *directory-name*

3. Enter the following command:

   `$SCHRODINGER/maestro &`

   The Maestro main window is displayed, and the working directory is Maestro's current working directory. If you are using an existing Maestro session, you can change the directory by choosing Change Directory from the Maestro menu, navigating to the appropriate directory and clicking OK.

Windows:

1. Double-click the Maestro icon on the desktop.

   You can also use the Start menu. Maestro is in the Schrödinger submenu.

2. From the Maestro menu, choose Change Directory.

3. Navigate to the working directory and click OK.

Receptor Grid Generation

This chapter contains exercises that demonstrate how to use the Receptor Grid Generation panel to set up and start a grid file calculation job. Grid files represent physical properties of a volume of the receptor (specifically the active site) that are searched when attempting to dock a ligand. You will use the grid files calculated in this chapter to dock ligands in later Glide exercises.

If you have not started Maestro, start it now. Before proceeding with the exercises, change the working directory to the `grids` directory. See Section 1.3 on page 3 for instructions on how to do these tasks. When you have finished the setup, you should see a path in the title bar of the main window that ends in grids.

2.1 Importing the Prepared Structures

The complex for this exercise is actually in two files, one containing the receptor and one containing the ligand.

1. Click the Import structures toolbar button.

   

   The Import panel is displayed.

2. From the Files of type menu, ensure that Maestro is chosen.

3. Click Options.

   The Import Options dialog box opens.

4. Ensure that Import all structures, Replace Workspace, and Fit to screen following import are all selected.

5. From the Include in Workspace option menu, ensure that First Imported Structure is chosen.

6. Click Close in the Import Options dialog box.

7. In the Import panel, navigate to the `structures` directory and select the file `1fjs_prep_recep.mae.gz`.

8. Click Open.

   The prepared protein is displayed in the Workspace. The protein structure is displayed in ribbon representation. The structure includes solvent molecules (glycerine) and ions (Ca^{2+} and Cl^{-}), but does not include the ligand.

9. Next import the ligand structure file by clicking the Import structures toolbar button.

   

   The Import panel is displayed.

10. From the Files of type menu, ensure that Maestro is chosen.
11. Click Options.

    The Import Options dialog box opens.

12. Deselect Replace Workspace and Fit to screen following import.
13. Click Close in the Import Options panel.
14. In the Import panel, navigate to the `structures` directory and select the file `1fjs_prep_lig.mae.gz`.
15. Click Open.

    The prepared ligand is displayed in the Workspace in tube representation.

2.2 Defining the Receptor

The receptor structure used for grid generation is taken from the Workspace, so you need to exclude the ligand atoms from consideration as part of the receptor.

1. From the Applications menu in the main window, choose Glide > Receptor Grid Generation.

   The Receptor Grid Generation panel opens with the Receptor tab displayed.

2. In the Define receptor section, ensure that Pick to identify ligand and Show markers are selected, and that, in the option menu, Molecule is chosen.
3. In the Workspace, pick an atom in the ligand molecule.

   Dark green markers appear on the ligand.

4. In the Van der Waals radii scaling section, ensure that Scaling factor is set to the default value of 1.00 (no scaling.)

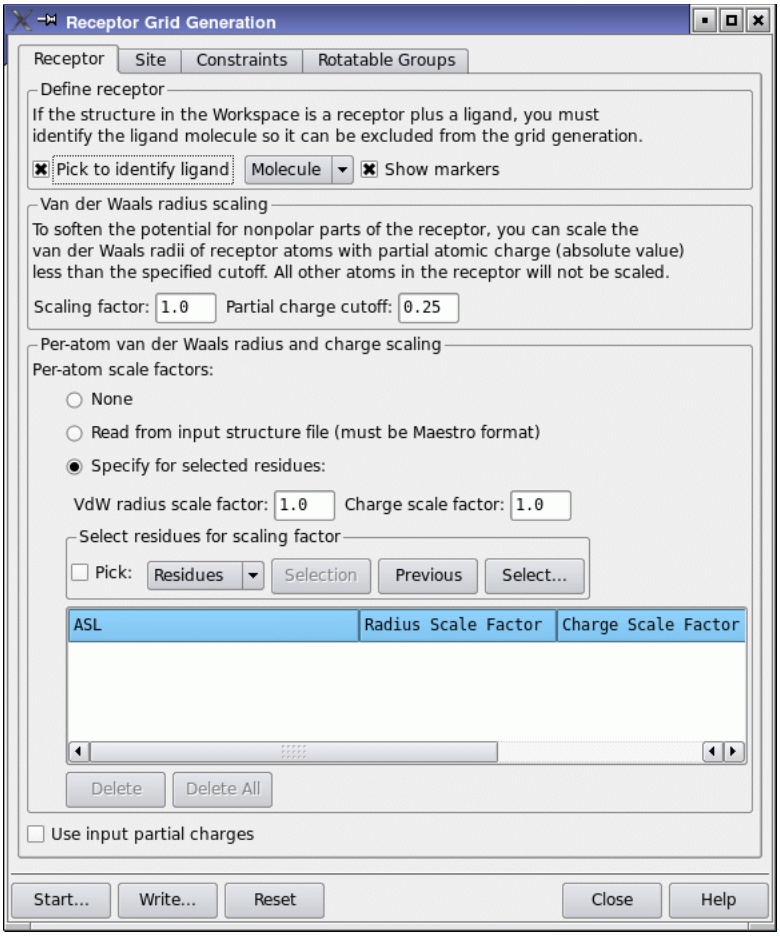

Figure 2.1. The Receptor Grid Generation panel.

2.3 Defining the Active Site

Now that the ligand has been excluded, the volume for which grids will be calculated can be defined:

1. Click the Site tab.

   The entire complex is shown with several types of markers: the dark green ligand molecule markers that appeared when the ligand was identified, and the new markers that appeared when the Site tab was opened:

   - The *enclosing box* is shown in purple.
   - The center of the enclosing box is marked by green coordinate axes.

The purple enclosing box represents the volume of the protein for which grids will be calculated. Generally, you should make the enclosing box as small as is consistent with the shape and character of the protein's active site and with the ligands you expect to dock.

2. In the Site tab, ensure that the Center option selected is Centroid of Workspace ligand.
3. Ensure that the Size option selected is the default, Dock ligands similar in size to the Workspace ligand.

   If you have a representative ligand in the active site, the default generates an enclosing box that is large enough for most systems. However, if you think that conformations of active ligands may exist that are significantly larger than the cocrystallized ligand, you should consider enlarging the enclosing box using the Dock ligands with length <= option.

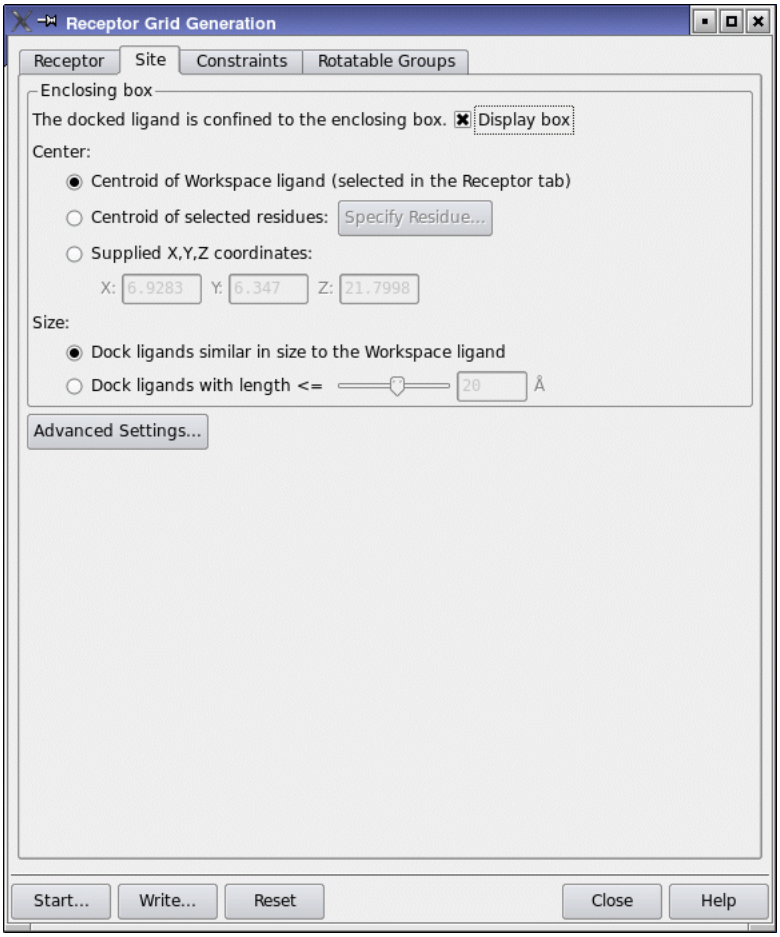

Figure 2.2. The Site tab of the Receptor Grid Generation panel.

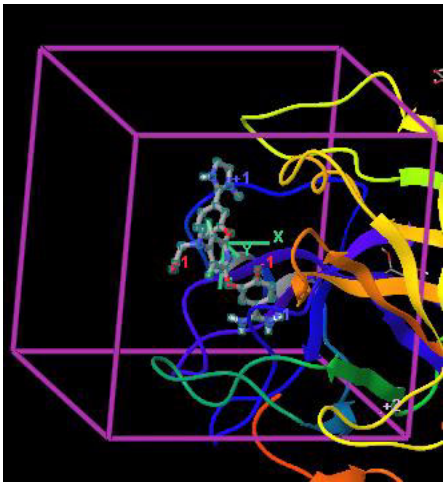

Figure 2.3. The marked ligand with enclosing box.

2.4 Setting Up Glide Constraints

The Constraints tab of the Receptor Grid Generation panel is used to define Glide constraints. In this exercise, you will define two constraints: a positional constraint and an H-bond constraint. To make it easier to see the parts of the receptor close to the ligand and to see the ligand atoms, you will first change the display.

For more information on using Glide constraints, see Section 4.4 of the *Glide User Manual*.

2.4.1 Setting the Display for Constraint Definition

1. From the Display only selected atoms toolbar button menu, choose Molecules, and click on a ligand atom.

   

   The ligand is displayed, and the receptor remains displayed as ribbons.

2. From the Show, hide, or color ribbons button menu, choose Delete Ribbons.

   

3. From Display residues within N Å of currently displayed atoms, choose 3 Å.

   

The residues that are closest to the ligand are displayed.

4. From the Undisplay toolbar button menu, choose Nonpolar hydrogens.

This action leaves the polar hydrogens displayed, making it easier to see the polar hydrogens that are likely to form hydrogen bonds.

2.4.2 Defining a Positional Constraint

1. In the Constraints tab of the Receptor Grid Generation panel, click the Positional tab.
2. Click New.

   The New Position dialog box opens.
3. In the Select atoms to define a position section, ensure that Pick is selected; and from the Pick option menu, that Atoms is selected.
4. Click the carbon atom that is between the two nitrogen atoms in the imidazole ring in the ligand.

   A semi-transparent gray sphere is displayed around the atom.
5. Enter the name `S4_arom` in the Name text box.
6. Enter `2.0` in the Radius text box.
7. Click OK.

   The constraint is added to the Positions table in the Positional tab, and the sphere changes to yellow. The name is displayed next to the sphere.

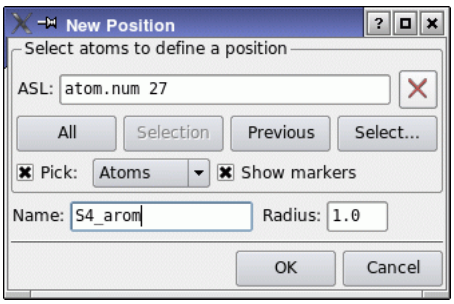

Figure 2.4. The New Position dialog box.

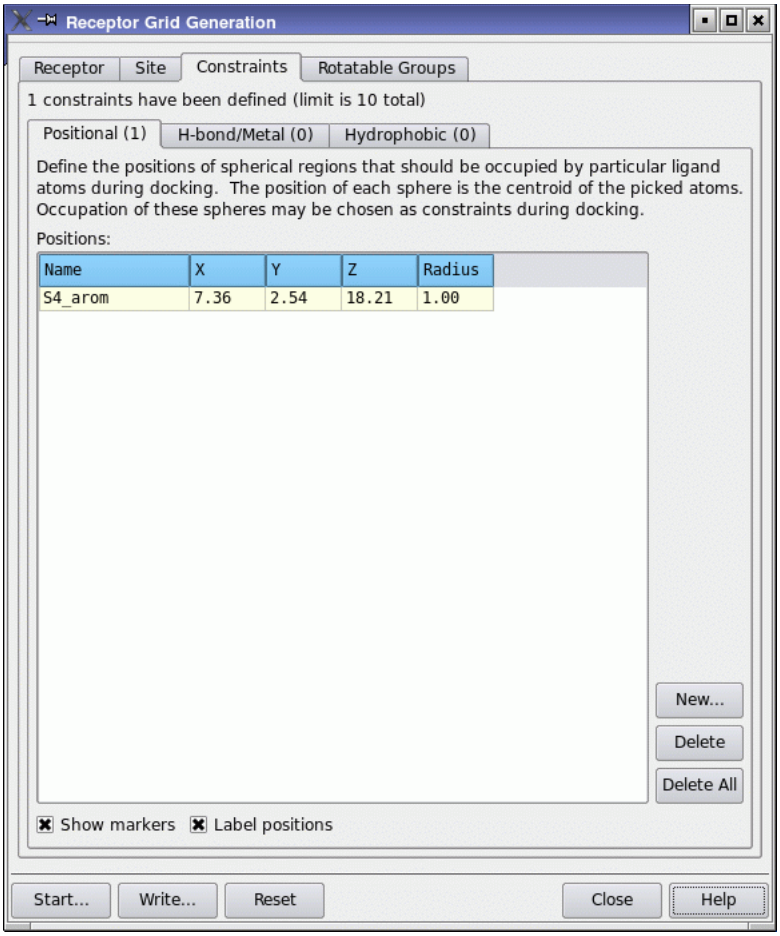

Figure 2.5. The Constraints tab of the Receptor Grid Generation panel showing the Positional subtab.

2.4.3 Defining an H-bond Constraint

Next, you will define an H-bond constraint, for the carboxylate that is H-bonded to the amidine of the ligand. To aid the picking of the constraint, H-bonds to the ligand will be displayed.

1. From the Display H-bonds button menu, choose Inter H-bonds, and click on a ligand atom.

   

   The hydrogen bonds between the ligand and the receptor appear as yellow dashed lines.

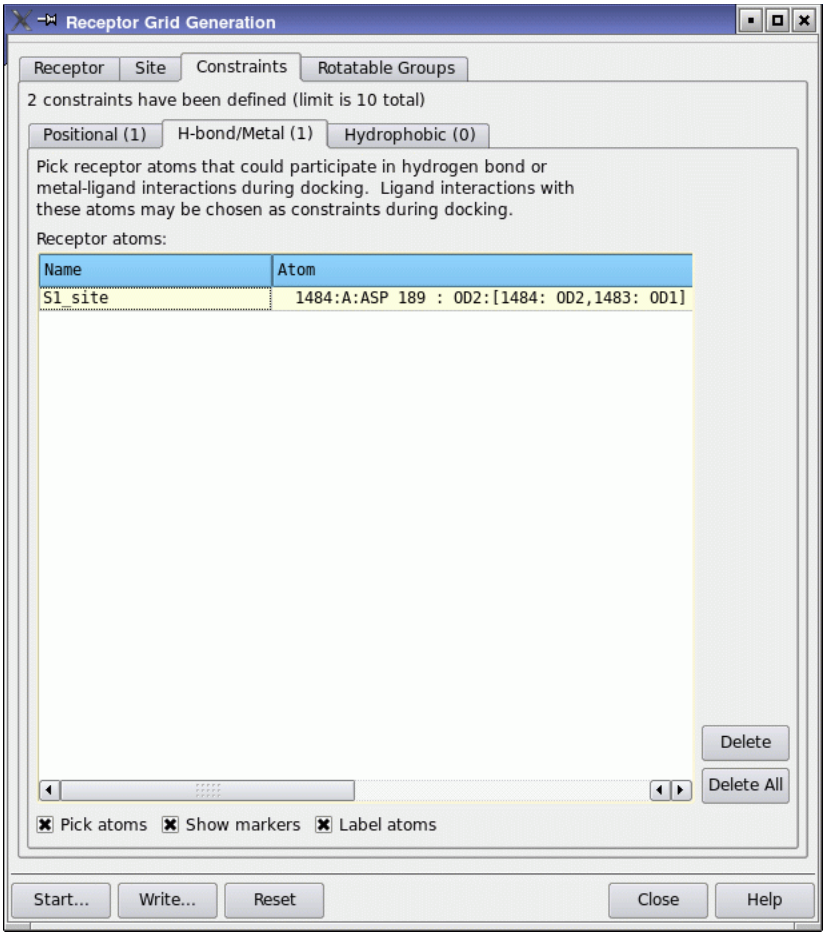

Figure 2.6. The Constraints tab of the Receptor Grid Generation panel showing the H-bond/Metal subtab.

2. In the Constraints tab, click on the H-bond/Metal tab.
3. Ensure that Pick atoms is selected.
4. Click on the carboxyl oxygen that is hydrogen-bonded to the amidine of the ligand.

   This atom is the OD1 atom of ASP 189. When you have picked the atom, an entry is added to the Receptor atoms table. The Name column shows the identity of the atom, with hbond in parentheses. In the Atom column, both oxygen atoms of the carboxylate are listed in square brackets, because Glide includes symmetry-related atoms as part of the constraint. Both atoms are marked in the Workspace with a padlock icon, if Show markers is selected, and the name is also displayed if Label atoms is selected.

5. Name the constraint `S1_site`, by editing the text in the Name column.
6. From the Display H-bonds button menu, choose Delete H-bonds.

2.5 Starting and Monitoring the Grid Calculation

With the ligand and the active site defined, and constraints set up, the grid generation job can be started.

1. In the Receptor Grid Generation panel, click Start.

   You might see a dialog box that warns about an unidentified ligand. This is a check for ligand-sized molecules that might be in the active site. It has found the glycerine molecules, which are not a problem for docking. You can therefore ignore the warning and click Continue.

   The Receptor Grid Generation - Start dialog box is displayed.

2. In the Output section, ensure that the Directory for grid files is the default, `./` , your current working directory, and ensure that Compress is selected.

   When this option is selected, the grid files are placed in an archive that is compressed, so it is easy to copy. Glide can make use of this compressed archive directly, and extracts the files from it as needed.

3. Check the main window title bar to confirm that the current working directory is *yourpath*`/tutorial/grids`.

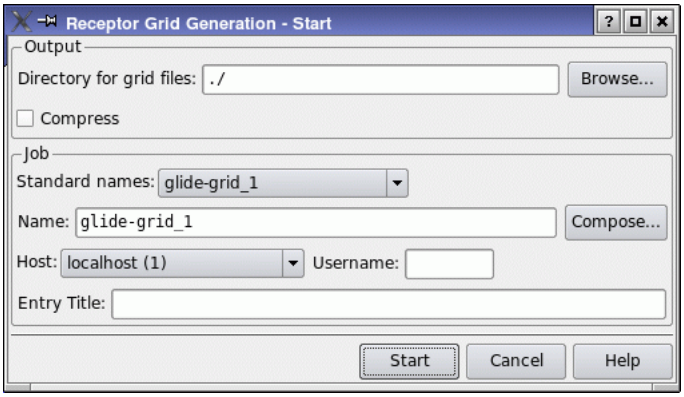

Figure 2.7. The Receptor Grid Generation - Start dialog box.

4. In the Job section, change the Name to `factorXa_grid`.
5. Choose a host and, if necessary, specify a user name.
6. Start the job by clicking Start.

   A warning dialog box might be displayed, informing you that the structure has not been prepared with the Protein Preparation Wizard. You can ignore this warning and click Continue.

   After a moment, the Monitor panel is displayed, and the job starts. While the job is in progress, the Status column displays "running." When the job is complete, the status is changed to "completed : finished". (The opening of the Monitor panel depends on a preference that is set in the Jobs tab of the Preferences panel.)

The job takes approximately 10 minutes on a 1 GHz Pentium 4 processor; this time may vary depending on your particular system configuration and workload.

Before the job is launched, these job input files are written:

`factorXa_grid.in`	Command input for grids job
`factorXa_grid.maegz`	Receptor structure input for grids job

When the calculation is complete, the `grids` directory will contain the following output files:

`factorXa_grid.log`	Log summary file from grids job
`factorXa_grid.out`	Output summary file from grids job
`factorXa_grid.zip`	Archive containing grid files

The zip archive contains the following grid files:

`factorXa_grid.csc`	`factorXa_grid.gsc`
`factorXa_grid.site`	`factorXa_grid.save`
`factorXa_grid_greedy.save`	`factorXa_grid.cons`
`factorXa_grid_recep.mae`	`factorXa_grid.grd`
`factorXa_grid_coul2.fld`	`factorXa_grid_vdw.fld`
`factorXa_grid.vdwc`	

In the next chapter, you will dock ligands using these grids. You can close the Receptor Grid Generation panel now.

Ligand Docking

The exercises in this chapter demonstrate the use of Glide to screen a multiple-ligand file for structures that interact favorably with a receptor active site. The receptor grid files you calculated in the previous chapter will be used to dock ligands from the file `50ligs.mae`. Most of the 50 ligands in the file are decoys, selected as a representative sample from a database of druglike molecules using the `ligparse` utility. Four ligands out of the total of 50 are active ligands for the chosen receptor. These ligands have all been prepared for docking with LigPrep—see Section 3.4 of the *Glide User Manual* or the *LigPrep User Manual* for more information on ligand preparation.

Typically, Glide standard-precision docking is used to find probable good binders in a large set; the top-scoring 10% to 30% can then be investigated more intensively using Glide extra-precision (XP) docking or other methods available from Schrödinger. In these exercises, you will use all three docking modes (HTVS, SP, and XP), and also investigate the use of constraints.

If you have not started Maestro, start it now (see Section 1.3).

Before proceeding with the exercises, change the working directory to the `glide` directory. See Section 2.1 on page 5 for instructions on how to do this.

3.1 Specifying a Set of Grid Files and Basic Options

In this exercise, you will select the grid that you calculated in Chapter 2 for the ligand docking job, and set the basic docking options.

1. Click the Clear workspace toolbar button.

   

2. From the Applications menu, choose Glide > Ligand Docking.

   The Ligand Docking panel opens with the Settings tab displayed.

3. In the Receptor grid section, click the Browse button.

   A file selector opens.

4. Navigate to the `tutorial/grids` directory, choose `factorXa_grid.zip`, and click Open.

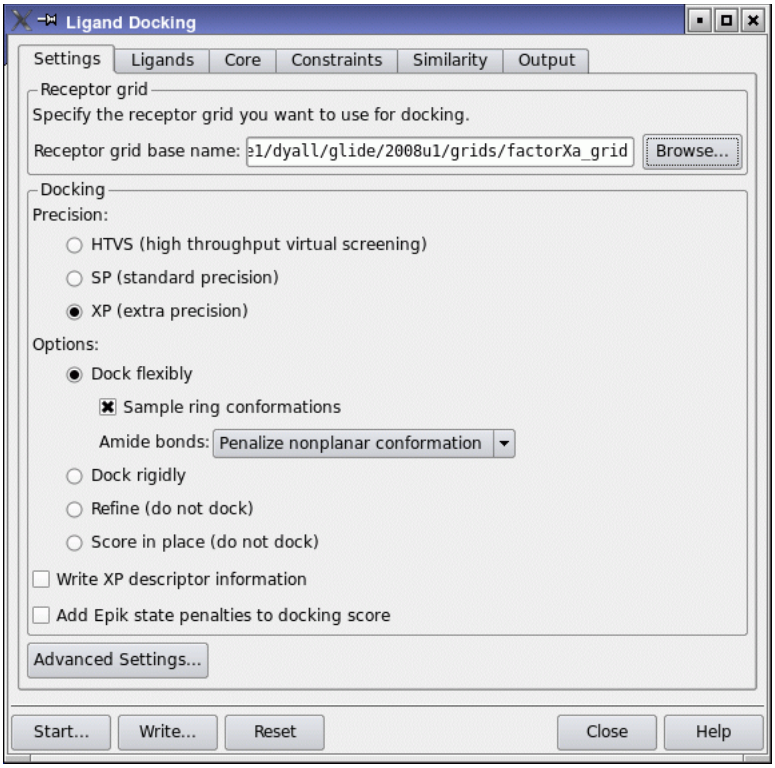

Figure 3.1. The Settings tab of the Ligand Docking panel.

The Receptor grid base name is *fullpath*`/tutorial/grids/factorXa_grid`.

5. In the Docking section, ensure that the Precision option is SP (standard precision).

   This is usually the best choice for docking large numbers of ligands. For more rapid screening you can use the HTVS (high throughput virtual screening) option. You will do this in a later exercise.

6. Under Options, ensure that Dock flexibly and Sample Ring Conformations are selected, and Penalize nonplanar conformation is chosen from the Amide bonds option menu. (These are the defaults.)

The receptor grids and the basic Glide settings for the ligand docking job are now specified. In the next section, you will specify a set of ligands to dock, and in the following section, you will specify output options. The options in the remaining three tabs, Core, Constraints, and Similarity, can be left at their defaults for this exercise.

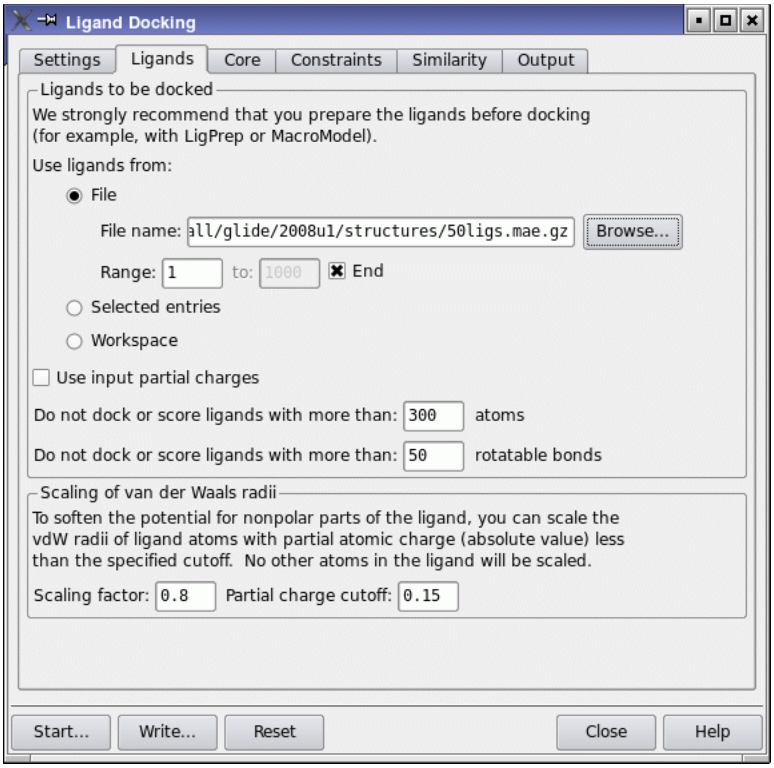

Figure 3.2. The Ligands tab of the Ligand Docking panel.

3.2 Specifying Ligands To Dock

There are several methods for specifying ligand structures to be docked with receptor grids. In this tutorial, you will specify a file containing a set of 50 ligands.

1. In the Ligands tab, ensure that File is selected.
2. Click Browse.

   A file selector is displayed. Ensure that Files of type is set to Maestro
3. Navigate to the `tutorial/structures` directory, choose `50ligs.mae.gz`, and click Open.
4. Ensure that the selected Range is from `1` to End (the default).
5. Ensure that van der Waals radii scaling for ligand atoms is set to the default values: Scaling factor to 0.80 and Partial charge cutoff to 0.15.

In this docking job the constraints that were specified in the grid generation will not be used. In a later exercise you will use the constraints to dock the same ligands.

3.3 Specifying Output Quantity and File Type

The Output tab allows you to specify the type of file to create for the output ligand poses and to determine how many poses to write, per ligand and per docking job.

1. In the Output tab, ensure that Write pose viewer file (includes receptor; filename will be <jobname>_pv.mae) is selected.

   You are specifying that the structural output from the docking job be written to a "pose viewer" file, a file of ligand poses that begins with the structure of the receptor. Having the receptor structure included in the file is convenient for displaying hydrogen bonds and contacts between the ligand and the receptor. In Chapter 4, you will use tools in the Project Table to examine the poses in the file `factorXa_sp_pv.maegz`.

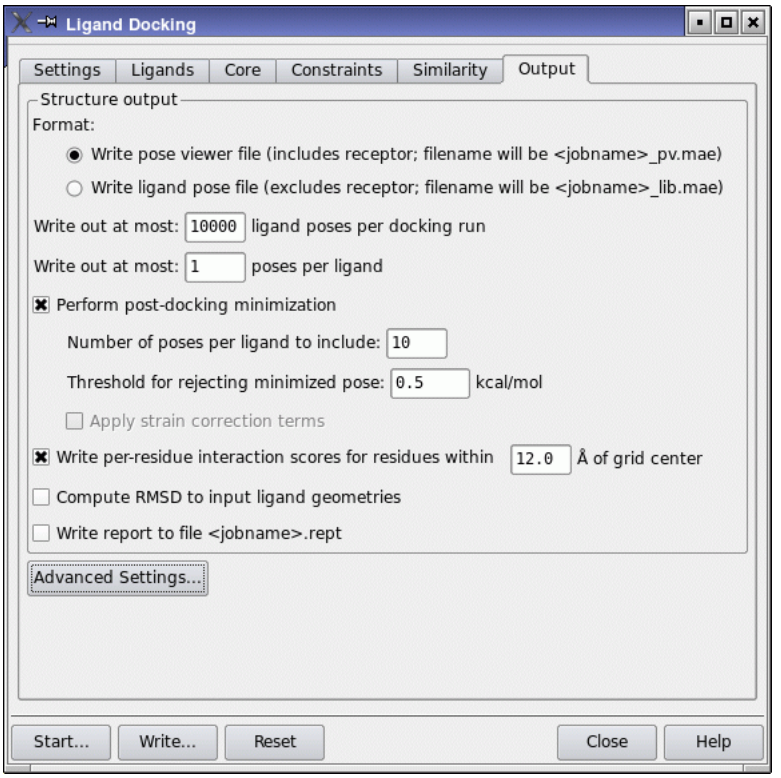

Figure 3.3. The Output tab of the Ligand Docking panel.

2. Ensure that the value of *m* in the Write out at most *m* poses per ligand text box is 1, the default.

   Because there are only 50 ligands in the input file, this setting ensures that no more than 50 poses, one for each ligand, will be collected and written to the pose viewer file.

3. Ensure that Perform post-docking minimization is selected, with the default number of poses per ligand (which is 5).

   Post-docking minimization in the field of the receptor produces better poses and only adds a small amount to the time taken.

3.4 Setting Up Distributed Processing

If you have access to a host machine with multiple CPUs, Glide can divide your multiple-ligand docking job into subjobs that can be distributed over several processors. The Start dialog box allows you to specify the number of subjobs and the number of processors to use. In this section, it is assumed that a host with five or more processors is available. If you have access to such a host, you can follow the instructions in this section without alteration. If the host has fewer than five processors, or you do not want to use five of its processors for this job, change the settings below accordingly.

If you cannot, or do not want to, distribute processing over multiple CPUs, skip to Section 3.5.

1. Click Start.

   The Start dialog box opens.

2. Change the value in the Separate docking into *N* subjobs text box to 10.

   The docking job will be split into 10 subjobs, each one docking 5 ligands. For the sake of speed, the number of ligands per subjob in this exercise is much smaller than would be typical in actual use. The default for *N* is 1, meaning that the job is run as a single job.

3. Choose a multiprocessor host from the Host list table.

   The number of processors on each host, as given in the hosts file, is shown in the Processors column.

4. Edit the value in the Use column for the selected host to specify the number of processors to use.

   The value in the Total to use text box is updated to reflect the value you entered. If you want to use multiple hosts, you can select multiple table rows, and specify the number of processors to use on each. The total number is given in the Total to use text box.

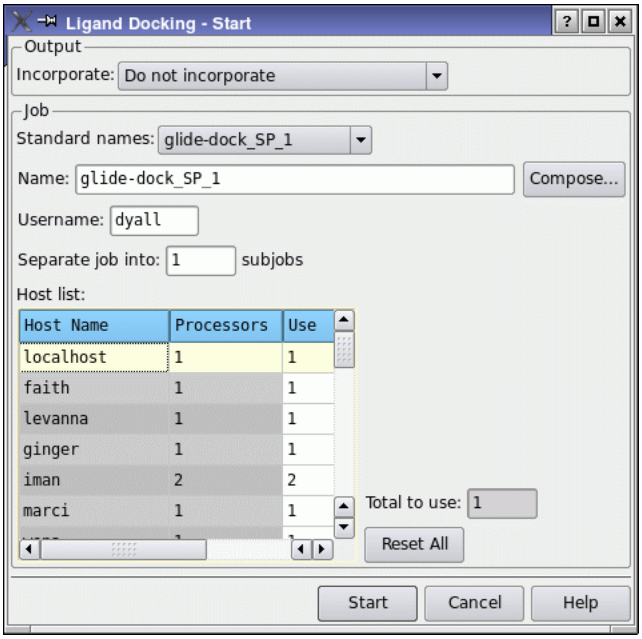

Figure 3.4. The Start dialog box, default settings.

The 10 subjobs will be distributed over the number of processors you specify. We suggest you use five processors. Then the first five subjobs will be sent first, followed by each of the remaining five as processors become available.

When you are satisfied with your distributed processing settings, continue to Section 3.5.

3.5 Starting the Ligand Docking Job

1. If you did not do so in the previous section, click Start.

   The Start dialog box opens.

2. In the Start dialog box, change the name to `factorXa_sp`.
3. From the Incorporate option menu, choose Append new entries as a new group.
4. If you have not already done so, select a host from the Host list.
5. If you have a different user name on the host you have selected from that on your local host, enter the correct user name in the Username text box.
6. Click Start.

   The docking job starts and the Monitor panel is displayed.

For the distributed processing example, as soon as the `factorXa_sp` job has been launched, it is divided into subjobs. As each subjob is launched on a processor, it is listed in the Monitor panel. When one subjob is finished, the next one is launched. To view the log for any subjob, select it in the job table and click Monitor. If the subjob is already finished, the entire log can be scrolled through in the File tab of the Monitor panel. The results for each subjob are stored in subdirectories of the output directory, and collected at the end into the output directory.

The time required for Glide docking jobs depends on the processor speed and workload, the size and flexibility of the ligands, and the volume specified by the enclosing box. As a rough estimate, docking a typical drug-like ligand takes about one minute on a 1.2 to 1.5 GHz Pentium 4 processor under Linux, using Glide 5.5 default settings. On a similar machine with a single processor, the 50-ligand docking job will usually finish in about 45 minutes. As 10 subjobs distributed over 5 similar processors, the docking job will finish in about 10 minutes.

When the job finishes, examine the results in the Project Table panel. The four active ligands (glide lignum 1 through 4) are ranked highest.

3.6 Docking in High-Throughput Virtual Screening Mode

Glide has a set of predetermined options that can speed up the docking by a factor of about seven over the standard precision (SP) docking mode. In this exercise, you will run an HTVS docking job on the same set of ligands as used in the SP docking exercise.

1. In the Settings tab, select HTVS (high throughput virtual screening).

   The other settings will be left as they were for the SP docking job.

2. In the Output tab, ensure that Write ligand pose file (excludes receptor; filename will be <jobname>_lib.mae) is selected.

3. Click Start.

   The Start dialog box opens. You should not need to run this job on multiple processors.

4. Change the job name to `factorXa_htvs`.

5. Select a host, and set the number of processors and subjobs to 1.

6. Click Start.

   The job should take only a few minutes to run.

When the job finishes, examine the results in the Project Table panel. The four active ligands are ranked in the top 5 ligands. Their scores differ a little from those in the SP docking run.

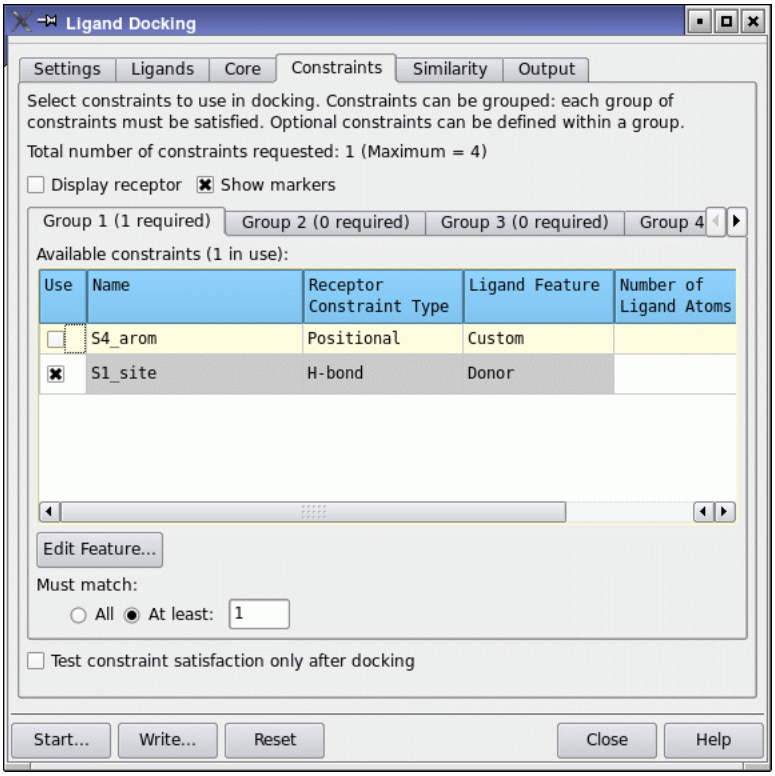

Figure 3.5. The Constraints tab of the Ligand Docking panel.

3.7 Docking Ligands Using Constraints

In this exercise, you will apply the constraints you defined in the grid generation to the docking of the same set of ligands as for the standard SP job. By default, no constraints are applied, even if they are defined. Here, you will require any one of the three constraints to be applied.

1. In the Settings tab, select SP (standard precision).

   The other settings will be left as they were for the HTVS docking job.

2. In the Constraints tab, click both check boxes in the Use column.

   These check boxes mark the constraint for use in docking.

3. Under Must match, select All.

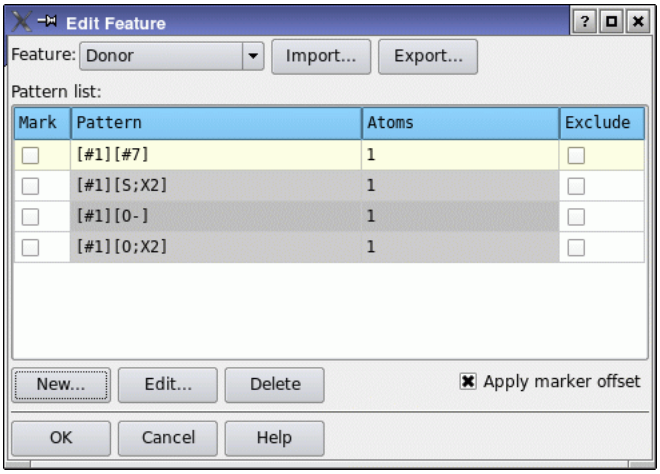

Figure 3.6. The Edit Feature dialog box.

For H-bond and hydrophobic constraints, the ligand features that must match these constraints are predefined. You can edit them if you want, but this is not necessary. For positional constraints, you must define the ligand feature that matches the constraint. Features are defined in terms of SMARTS patterns.

4. From the Available constraints table, select the S4_arom row and click Edit Feature.

   The Edit Feature dialog box opens. There are no SMARTS patterns in the Pattern list table, because the Custom feature type is undefined by default.

5. Click New.

   The New Pattern dialog box opens.

6. Enter the following text into the SMARTS pattern text box:

   `[a]`

   This pattern matches all aromatic atoms.

7. Enter `1` into the Numbers text box.

   This is the index of the atom in the SMARTS pattern that is matched by the constraint. Since there is only one atom in the pattern, this is the only possible choice.

8. Click OK.

   The New Pattern dialog box closes, and a row is added to the Pattern list table in the Edit Feature dialog box.

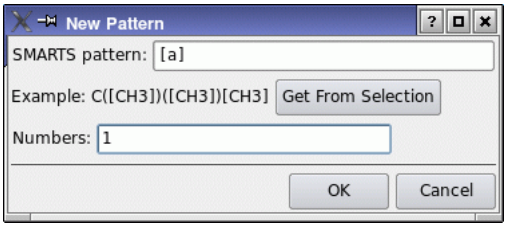

Figure 3.7. The New Pattern dialog box.

9. Click OK.

   The Edit Feature dialog box closes. This completes the definition of the Custom feature. If you did not define this feature, the docking job would not be started.

10. Click Start.

    The Start dialog box opens. If you ran the SP docking job on multiple processors before, you can do the same for this job.

11. Select the host and number of processors.

12. Change the job name to `factorXa_sp_cons` and click Start.

    The Monitor panel opens and displays the progress of the job.

When the job finishes, examine the results in the Project Table. Of the 50 ligands, poses are reported for only six ligands, of which the first four are the actives, and the other two did not score very well. The remaining ligands did not satisfy the constraints. These results indicate that the application of constraints serves to discriminate between ligands that bind in the proper mode, and ligands that don't.

3.8 Docking Ligands Using Core Constraints

In this exercise, you will apply core constraints defined by a pattern in the reference ligand, and use these to dock a set of structures. Structures that do not include the core pattern will not be docked. The first task is to import the reference ligand.

1. In the main window, on the main toolbar, click the Import structures button.

   

   The Import panel is displayed.

2. From the Files of type menu, ensure that Maestro is chosen.

3. Click Options.

   The Import Options dialog box opens.

4. Ensure that Import all structures, Replace Workspace, and Fit to screen following import are all selected.

5. From the Include in Workspace option menu, ensure that First Imported Structure is chosen.

6. Click Close in the Import Options dialog box.

7. Navigate to the `structures` subdirectory and select the file `sar_reference.mae.gz`.

8. Click Open.

   The reference ligand is displayed in the Workspace.

Next, the settings for the previous constraints job need to be cleared.

9. In the Ligand Docking panel, click the Settings tab, and select SP (standard precision).

   The other settings will be left as they were for the previous docking job.

10. In the Constraints tab, clear all check boxes in the Use column.

    Receptor-based constraints will no longer be applied; instead you will be using ligand-based constraints.

The core constraint is defined in the following steps.

11. In the Core tab, select Restrict docking to reference position.

    The first few controls in the Define core section become available.

12. Enter `1.5` in the Tolerance text box.

13. Click an atom in the reference ligand (in the Workspace).

    The ligand is marked with purple markers, and the remaining controls in the Define core section become available.

14. Under Core atoms, select SMARTS pattern.

15. In the main window, from the Undisplay toolbar button menu, choose Nonpolar hydrogens.

    

16. Rotate the structure so that you can clearly see the three six-membered rings.

17. Ensure that the Workspace selection button is selected (indented) and displays an A (for picking atoms).

18. Select the three six-membered rings with their ether linkages, and the amidine group on the terminal ring.

    Do not include the carboxyl on the middle ring or the hydroxyl on the terminal ring.

    You can drag to make the first selection, then hold down the SHIFT key and drag or click to add atoms to the selection. The selected atoms are marked in yellow rather than in purple.

19. In the Ligand Docking panel, click Get From Selection.

    The Smarts pattern text box is filled in with the pattern that corresponds to the atoms selected in the Workspace. The markers on the Workspace selection turn green.

Finally, the ligands to be docked need to be selected. A different set is used from the set used for the previous runs.

20. In the Ligands tab, ensure that File is selected.
21. Click Browse.

    A file selector is displayed. Ensure that Files of type is set to Maestro.

22. Navigate to the `tutorial/structures` directory, choose `sar_series.mae.gz`, and click Open.
23. Ensure that the selected Range is from `1` to End (the default).
24. Start the job with the name `factorXa_sp_core`.

    The job takes a similar time to the constraints job.

25. When the job finishes, examine the results in the Project Table.

3.9 Refining Docked Ligands with Glide XP

In this exercise, you will use Glide XP to refine a set of ligands taken from the first SP docking run.

1. In the Core tab, select Do not use.

   Core constraints are now turned off.

2. In the Settings tab, under Precision, select XP (extra precision).
3. Under Options, select Refine (do not dock).
4. Select Write XP descriptor information.

   Note: This option requires a special license. If you do not have this license, do not select the option. You will obtain results without the license, but you will not be able to complete the exercise in Section 4.5 on page 33.

5. In the Ligands tab, ensure that File is selected.
6. Click Browse.

   A file selector is displayed.

7. Ensure that Files of type is set to Maestro.
8. Navigate to the `tutorial/structures` directory.
9. Choose `refine_xp_entries.mae.gz`, and click Open.
10. Ensure that the selected Range is from `1` to End (the default).
11. If you did not select Write XP descriptor information, in the Structure output section of the Output tab, select Write pose viewer file (includes receptor; filename will be <job-name>_pv.mae.gz).

    If you did select this option, the Write pose viewer file option is automatically selected.

12. Start the job with the name `factorXa_xp_refine`.

This job may take up to 30 minutes to run on a 2 GHz processor.

Examining Glide Data

In this chapter, Glide results are examined with an emphasis on visual rather than numerical appraisal. The first set of exercises use the Project Table to display the results of a Glide docking job, examine individual ligand poses and their contacts with the input receptor structure. The second set of exercises uses the Glide XP Visualizer panel to display information on the terms in the Glide XP scoring function that contribute to the ligand binding.

If you have not started Maestro, start it now (see Section 1.3).

Before proceeding with the exercises, change the working directory to the `glide` directory. See Section 1.3 on page 3 for instructions on how to do this.

4.1 Importing Pose Data

The first task is to import poses from a pose viewer file into the Maestro project.

1. In the main window, on the toolbar, click the Import structures button.

   

   The Import panel opens.

2. Click on Options.

   The Import Options dialog box opens.

3. Ensure that Maestro is chosen from the Files of type menu.

4. Ensure that Import all structures, Replace Workspace, and Fit to screen following import are all selected.

5. Click Close in the Import Options dialog box.

6. In the Import panel, select the file `factorXa_sp_pv.maegz` and click Open.

The receptor and the ligands are imported as an entry group named factorXa_sp_pv. The receptor is displayed in the Workspace.

4.2 Viewing Poses

The Project Table panel has special options for entry groups that consist of a receptor and a set of ligands. These options are available from the Entry menu, under View Poses, when you have a single entry group selected.

1. Open the Project Table panel.

   You can do this by clicking the Open/Close project table toolbar button.

   

   The entries in the entry group containing the receptor and the poses that you imported should be selected. If not, click in the Row column for the entry group.

2. From the Entry menu, choose View Poses > Setup.

   The receptor is fixed in the Workspace, and the first pose is included in the Workspace. A Mark property is added to the Project Table so you can record any poses that you want to mark as being of special interest. To mark a Workspace entry, type M.

The next few steps change the display so that the ligand fills most of the Workspace.

3. From the Workspace selection toolbar button, choose Molecule.

   

   The A on the button changes to an M, to indicate that molecules are being picked.

4. Click the ligand molecule.

   The atoms in the ligand are marked with pale yellow markers.

5. Click the Fit to screen toolbar button.

   

   The view zooms in so that the ligand fills most of the Workspace.

6. Click in an empty portion of the Workspace to clear the selection.

7. Next turn on Workspace Feedback by going to Display > Show Single-Entry Feedback or by typing S.

By default Workspace Feedback appears in the upper right corner of the Workspace, and includes the entry Title, GlideScore, and EmodelScore.

Now you are ready to view the poses.

8. Press the RIGHT ARROW key.

   The second pose replaces the first in the Workspace. The RIGHT ARROW and LEFT ARROW keys can be used to step through the selected poses.

9. Shift-click the entry for the first pose.

   The first pose is added to the Workspace.

10. Press the RIGHT ARROW key.

    The third pose replaces the first two in the Workspace.

11. Reselect the first pose by clicking its In column.

4.3 Displaying Atoms by Proximity

In this section, you will select a display that includes the ligand and the receptor residues nearest the ligand. This is useful for examining contacts and hydrogen bonds between the ligand and the active site of the receptor.

1. From the Display only selected atoms button menu, choose Molecules.

   

2. Click on an atom in the ligand.

   The ligand molecule is displayed, and all other atoms are undisplayed.

3. From the Display residues within N Å of currently displayed atoms button menu, choose +5 Å.

   

   Only the ligand and the nearby residues are displayed. All residues that do not have any atoms within 5 Å of the ligand are undisplayed. Hiding the residues that do not come into contact with the ligand makes it easier to examine the ligand-receptor interactions.

   You can also open the Atom Selection dialog box from the Display only button menu to pick the ligand and add atoms within a given radius of a particular set of atoms. This approach allows you more flexibility in picking the atoms.

4.4 Visualizing Hydrogen Bonds and Contacts

In this exercise, you will display hydrogen bonds between the ligand and the receptor. (To display hydrogen bonds between any two sets of atoms, use the Atom set 1 and Atom set 2 selection options in the H-Bonds tab.)

1. In the Project Table, choose Entry > View Poses > Display H-bonds.

   Hydrogen bonds to the currently displayed pose are displayed as yellow dashed lines.

   If you want to change the cutoffs for defining hydrogen bonds, choose Entry > View Poses > Define H-bonds, and change the values in the Measurements panel, which is displayed by choosing this item.

   If you want a count of hydrogen bonds between the ligands and the receptor, choose Entry > View Poses > Count H-bonds. An HBond property will be added to the Project Table, and the count may take a few seconds to finish.

2. Use the RIGHT ARROW and LEFT ARROW keys to step through the poses.

   The hydrogen bonds are displayed as each pose is included in the Workspace. Note the difference in hydrogen bonding patterns between the ligands.

In the pose list, 344 Good vdW, 6 Bad vdW, and 1 Ugly vdW contacts are reported for `1dwd`. Even the least-good ligand pose has many more good contacts than bad or ugly ones. The default is to display only Bad or Ugly contacts between the ligand and the receptor. (To display contacts between any two sets of atoms, use the Atom set 1 and Atom set 2 selection options in the Contacts tab.)

3. From the Entry menu, choose View Poses > Display Contacts.

   The contacts are shown as dashed lines connecting Workspace atoms. Ugly contacts are shown in red and Bad contacts are shown in orange. By default, atoms that are hydrogen bonded are not considered to have bad or ugly contacts.

   If you want to change the cutoffs to redefine the distance criteria for Good, Bad, and Ugly contacts, choose Entry > View Poses > Define Contacts, and change the values in the Measurements panel, which is displayed by choosing this item.

   If you want a count of contacts between the ligands and the receptor, choose Entry > View Poses > Count Contacts. The count may take a few seconds to finish, and three new properties are added to the Project Table: Good, Bad, and Ugly.

4. Use the RIGHT ARROW and LEFT ARROW keys to step through the poses.

4.5 Visualizing Glide XP Descriptors

In this exercise, you will use the Glide XP Visualizer to examine the contributions of various terms to the XP scoring function. The terms are given a spatial representation that you can display together with the ligand and the receptor.

Note: This exercise uses the results obtained from the exercise in Section 3.9 on page 26, which requires a special license for XP descriptor generation.

1. Click the Clear Workspace toolbar button.

   

2. From the Applications menu, choose Glide > XP Visualizer.

   The Glide XP Visualizer panel opens.

3. Click Open.

4. Select `factorXa_xp_refine.xpdes` in the file selector that opens, and click Open.

   After a short delay, the receptor and the highest-scoring ligand are displayed in the Workspace, and the table in the panel is filled in. The underlined values indicate that there is a corresponding visualization for this value.

5. Click the PhobEn cell for ligand 16088.

   You might want to deselect Narrow columns to see the entire column heading. This column displays the hydrophobic enclosure rewards. After a few seconds, the naphthalene of the ligand is displayed in ball-and-stick, and the hydrophobic atoms on the protein surrounding this ring are displayed in CPK in gray.

6. Click the PhobEn cell for ligand 612278.

   Note that for this ligand there is only a benzene ring rather than a naphthalene.

7. Click the HBond cell for ligand 16088.

   The hydrogen bonds are displayed as yellow dotted lines and annotated with their lengths.

8. Click the Electro cell for ligand 16088.

   The amidine nitrogens are displayed in ball-and-stick, to indicate their contribution to electrostatic rewards.

9. Click the π Stack cell for ligand 16088.

   The aromatic groups in the protein that contribute to the π stacking reward are displayed in gray in CPK, and the ligand aromatic group is displayed in ball-and-stick.

10. Click the RotPenal cell for ligand 16088.

    The bonds that contribute to the rotatable bond penalty are displayed as tubes.

11. Close the Glide XP Visualizer panel.

12. Clear the Workspace, deleting the scratch entry.

4.6 Finishing the Exercises

Close the scratch project you are working in. Because you have written the output structure files to your directory tree, you do not need to save the scratch project or Workspace structures. Click OK to delete any scratch entries.

From the Maestro menu, choose Quit and click Quit, do not save log file. (For more information about Quit panel options and `maestrolog.cmd` files, click Help instead.)

Getting Help

Schrödinger software is distributed with documentation in PDF format. If the documentation is not installed in `$SCHRODINGER/docs` on a computer that you have access to, you should install it or ask your system administrator to install it.

For help installing and setting up licenses for Schrödinger software and installing documentation, see the *Installation Guide*. For information on running jobs, see the *Job Control Guide*.

Maestro has automatic, context-sensitive help (Auto-Help and Balloon Help, or tooltips), and an online help system. To get help, follow the steps below.

- Check the Auto-Help text box, which is located at the foot of the main window. If help is available for the task you are performing, it is automatically displayed there. Auto-Help contains a single line of information. For more detailed information, use the online help.
- If you want information about a GUI element, such as a button or option, there may be Balloon Help for the item. Pause the cursor over the element. If the Balloon Help does not appear, check that Show Balloon Help is selected in the Maestro menu of the main window. If there is Balloon Help for the element, it appears within a few seconds.
- For information about a panel or the tab that is displayed in a panel, click the Help button in the panel, or press F1. The help topic is displayed in your browser.
- For other information in the online help, open the default help topic by choosing Online Help from the Help menu on the main menu bar or by pressing CTRL+H. This topic is displayed in your browser. You can navigate to topics in the navigation bar.

The Help menu also provides access to the manuals (including a full text search), the FAQ pages, the New Features pages, and several other topics.

If you do not find the information you need in the Maestro help system, check the following sources:

- *Maestro User Manual*, for detailed information on using Maestro
- *Maestro Command Reference Manual*, for information on Maestro commands
- *Maestro Overview*, for an overview of the main features of Maestro
- *Maestro Tutorial*, for a tutorial introduction to basic Maestro features
- *Glide User Manual*, for detailed information on using Glide
- *Protein Preparation Guide*, for information on protein preparation for Glide
- *Impact Command Reference Manual*, for information on Impact commands

- Glide Frequently Asked Questions pages, at https://www.schrodinger.com/Glide_FAQ.html
- Known Issues pages, available on the Support Center.

The manuals are also available in PDF format from the Schrödinger Support Center. Local copies of the FAQs and Known Issues pages can be viewed by opening the file `Suite_2009_Index.html`, which is in the `docs` directory of the software installation, and following the links to the relevant index pages.

Information on available scripts can be found on the Script Center. Information on available software updates can be obtained by choosing Check for Updates from the Maestro menu.

If you have questions that are not answered from any of the above sources, contact Schrödinger using the information below.

E-mail: help@schrodinger.com
USPS: Schrödinger, 101 SW Main Street, Suite 1300, Portland, OR 97204
Phone: (503) 299-1150
Fax: (503) 299-4532
WWW: http://www.schrodinger.com
FTP: ftp://ftp.schrodinger.com

Generally, e-mail correspondence is best because you can send machine output, if necessary. When sending e-mail messages, please include the following information:

- All relevant user input and machine output
- Glide purchaser (company, research institution, or individual)
- Primary Glide user
- Computer platform type
- Operating system with version number
- Glide version number
- Maestro version number
- mmshare version number

On UNIX you can obtain the machine and system information listed above by entering the following command at a shell prompt:

```
$SCHRODINGER/utilities/postmortem
```

This command generates a file named *username*`-`*host*`-schrodinger.tar.gz`, which you should send to help@schrodinger.com. If you have a job that failed, enter the following command:

```
$SCHRODINGER/utilities/postmortem jobid
```

where *jobid* is the job ID of the failed job, which you can find in the Monitor panel. This command archives job information as well as the machine and system information, and includes input and output files (but not structure files). If you have sensitive data in the job launch directory, you should move those files to another location first. The archive is named *jobid*`-archive.tar.gz`, and should be sent to help@schrodinger.com instead.

If Maestro fails, an error report that contains the relevant information is written to the current working directory. The report is named `maestro_error.txt`, and should be sent to help@schrodinger.com. A message giving the location of this file is written to the terminal window.

More information on the `postmortem` command can be found in Appendix A of the *Job Control Guide*.

On Windows, machine and system information is stored on your desktop in the file `schrodinger_machid.txt`. If you have installed software versions for more than one release, there will be multiple copies of this file, named `schrodinger_machid-`*N*`.txt`, where *N* is a number. In this case you should check that you send the correct version of the file (which will usually be the latest version).

If Maestro fails to start, send email to help@schrodinger.com describing the circumstances, and attach the file `maestro_error.txt`. If Maestro fails after startup, attach this file and the file `maestro.EXE.dmp`. These files can be found in the following directory:

```
%USERPROFILE%\Local Settings\Application Data\Schrodinger\appcrash
```

Glossary

Base Name—The name entered in the Base name for grid files text box that is used to write grid files during a grid file calculation, or to find pre-existing grid files during a docking job.

Bounding Box—The green, cube-shaped marker that appears in the Workspace during Glide docking job setup after you select active site residues, coordinates, or a ligand to be used as the box's center. The box represents the space in which ligands are allowed to move during docking. Increasing the size of the bounding box increases the space that can be sampled by the docked ligands, and consequently increases the CPU time required for the calculation.

Contacts—Graphical representations of the van der Waals interactions between the atoms of two or more molecules. Within Maestro, contacts are categorized as "Good," "Bad," and "Ugly." Good contacts are those that have van der Waals radii consistent with the experimentally determined values for the involved atom types. Bad contacts depict those interactions that are experimentally improbable. Ugly contacts represent van der Waals interactions that are disallowed in experimental systems.

Enclosing Box—The purple, cube-shaped marker that appears in the Workspace after you specify active residue sites, coordinates, or a ligand to be used as a bounding box center using the Glide panel. The enclosing box represents the space that any part of any specified ligand can sample during a docking calculation. Compare this with the green *bounding box*, which represents the space that the center of each specified ligand must be confined to during a docking calculation.

Flexible Docking—A job type in which alternate conformations for each ligand are generated during the docking process, and then the interactions between the receptor and the conformers are analyzed. After docking jobs are complete, the conformers, or "poses," are ranked according to their overall interaction with the receptor. The results can be posted to a pose view file, which can be examined using the Glide Pose Viewer panel.

GlideScore—Glide's scoring function (based on ChemScore). GlideScore is used in ranking ligand poses found in docking. In Liaison, GlideScore is used in an alternative binding energy model.

Grid Files—Files written by Glide during grid setup. These files contain data about the properties of the associated receptor and are used during docking.

Ligand Centroid—Used to define the enclosing box center, a ligand centroid is the point whose x, y, and z coordinates are the mean of the minimum and maximum *x*, *y*, and *z* coordinates of all the atoms in the ligand.

Pose Viewer Panel—An analysis tool that displays the results of Glide docking jobs. These results, which are recorded in a pose view file, include the ligand name, pose number, overall score, number of contacts, and other data. The poses within the file are arranged in the list according to score: ligands with the most energetically favorable interactions with the receptor appear at the beginning, and ligands with less favorable interactions appear near the end. The Glide Pose Viewer panel can also be used to visualize contacts and hydrogen bonds between ligand and receptor molecules, or to write structure files containing one or more ligand poses.

Reference Ligand—A user-specified structure whose ligand/receptor docking score will be compared with all other docked ligands.

Rigid Docking—A job type in which only supplied conformations of the specified ligands will be docked, scored, and displayed in a pose view file. This job type is useful if you have already performed a conformational search on the ligands that you want to dock.

120 West 45th Street, 29th Floor
New York, NY 10036

101 SW Main Street, Suite 1300
Portland, OR 97204

8910 University Center Lane, Suite 270
San Diego, CA 92122

Zeppelinstraße 13
81669 München, Germany

Dynamostraße 13
68165 Mannheim, Germany

Quatro House, Frimley Road
Camberley GU16 7ER, United Kingdom

SCHRÖDINGER®